

000 001 002 003 004 005 DIFFUSION NEGATIVE PREFERENCE OPTIMIZATION 006 MADE SIMPLE 007 008 009

010 **Anonymous authors**
011 Paper under double-blind review
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027

028 ABSTRACT 029

030 Classifier-Free Guidance (CFG) improves diffusion sampling by encouraging con-
031 ditional generations while discouraging unconditional ones. Existing preference
032 alignment methods, however, focus only on positive preference pairs, limiting
033 their ability to actively suppress undesirable outputs. Diffusion Negative Prefer-
034 ence Optimization (Diff-NPO) approaches this limitation by introducing a sep-
035 arate negative model trained with inverted labels, allowing it to capture signals
036 for suppressing undesirable generations. However, this design comes with two
037 key drawbacks. First, maintaining two distinct models throughout training and
038 inference substantially increases computational cost, making the approach less
039 practical. Second, at inference time, Diff-NPO relies on weight merging between
040 the positive and negative models, a process that dilutes the learned negative align-
041 ment and undermines its effectiveness. To overcome these issues, we introduce
042 Diff-SNPO, a single-network framework that jointly learns from both positive
043 and negative preferences. Our method employs a bounded preference objective to
044 prevent winner-likelihood collapse, ensuring stable optimization. Diff-SNPO de-
045 livered strong alignment performance with significantly lower computational over-
046 head, showing that explicit negative preference modeling can be simple, stable,
047 and efficient within a unified diffusion framework. Code will be released.
048

049 1 INTRODUCTION 050

051 Diffusion models (Ho et al., 2020) have become the backbone of modern visual content generation,
052 achieving remarkable fidelity in synthesizing images, videos, and multimodal content (Rombach
053 et al., 2022; Ho et al., 2022; Ruan et al., 2023). However, models trained on vast, uncurated web-
054 scale datasets often inherit biases and fail to align with human notions of quality, aesthetics, or
055 safety. Consequently, fine-tuning models with human feedback through preference alignment has
056 become a critical step for bridging the gap between a model’s raw capabilities and user intent (Black
057 et al., 2024; Wallace et al., 2024; Li et al., 2024).
058

059 At the heart of high-quality diffusion sampling lies *Classifier-Free Guidance (CFG)* (Ho & Sal-
060 mans, 2021), a technique that enhances sample quality by amplifying the contrast between con-
061 ditional and unconditional (or negatively conditioned) likelihoods. This mechanism guides generation
062 toward prompt-aligned outputs by explicitly pushing samples away from the lower-quality uncon-
063 ditional distribution. However, a challenge arises with many preference optimization methods, such
064 as Diffusion Direct Preference Optimization (Diff-DPO). These methods often apply the same opti-
065 mization objective to both the conditional and unconditional branches. This uniform reinforcement
066 of preferred attributes, while shifting the overall distribution in the desired direction, fails to heighten
067 the critical contrast that CFG relies on. Consequently, the model’s ability to generate outputs aligned
068 with user intent remains limited.
069

070 To further strengthen the contrastive effect of CFG, recent methods (Fu et al., 2025; Wang et al.,
071 2025) have incorporated negative preference alignment into the training process. In particular,
072 CHATS (Fu et al., 2025) and Diff-NPO (Wang et al., 2025) train two separate models: a “posi-
073 tive” model on standard preference data and a “negative” model on inverted preferences. This setup
074 allows the system to explicitly learn from dispreferred samples and steer generation away from un-
075 desirable attributes, resulting in improved human preference alignment. Despite their effectiveness,
076 a key limitation of both approaches lies in their reliance on distinct sets of parameter for the two
077

054 Table 1: Comparison of methods by alignment type, model setup, and use of merging strategy.
055

Method	Negative Alignment	Dual Model	Merging Strategy
DPO	✗	✗	✗
NPO	✓	✓	✓
CHATS	✓	✓	✓
SNPO	✓	✗	✗

062 models, which often leads to misaligned outputs and complicates guided sampling (Wang et al.,
063 2025). Diff-NPO addresses this by merging the weights of the two models, whereas CHATS per-
064 turbs conditional embeddings to blend their signals. These differences are summarized in Table 1. In
065 this work, we focus on Diff-NPO, since its sampling procedure more closely follows standard CFG.
066 Building on this foundation, we highlight two central limitations in its design. First, training and
067 sampling from two independent models inherently doubles both computational and memory costs,
068 creating a significant scalability challenge. Second, merging weights that were optimized separately
069 introduces a distribution mismatch that weakens the contribution of the negative model. This re-
070 duces the method’s ability to suppress dispreferred outputs and ultimately limits the gains expected
071 from negative alignment.

072 In this work, we ask: *Can we achieve the benefits of explicit negative modeling without the costs*
073 *and compromises of a dual-model architecture?* We introduce *Diffusion Simple Negative Prefer-
074 ence Optimization* (*Diffusion-SNPO*), a framework that integrates positive and negative preference
075 signals into a single, unified network. Our approach leverages the inherent dual-branch structure of
076 CFG-enabled models, training the conditional branch on positive preferences and the unconditional
077 branch on negative (inverted) preferences. This design eliminates the need for separate models and
078 weight merging, preserving a strong, explicit contrast between preferred and dispreferred distribu-
079 tions within one efficient architecture.

080 However, we find that naively applying this strategy results in an unintended side effect: the gener-
081 ated images become progressively blurrier as training continues. We attribute this to the likelihood
082 decrease observed in the DPO algorithm (Rafailov et al., 2024b; Pal et al., 2024; Cho et al., 2025),
083 which, when coupled with conflicting gradients from flipped preferences between the conditional
084 and unconditional branches, causes instability. As a result, the model converges to a suboptimal so-
085 lution, with blurring becoming more pronounced as the likelihood of winning samples decreases. To
086 address this, we adapt Bounded DPO (Cho et al., 2025), a preference optimization method designed
087 to increase the likelihood of winning samples during training, to Diffusion Models. This adaptation
088 stabilizes training by preventing the loss from being dominated by low-likelihood “losing” samples,
089 thereby boosting the likelihood of winning samples throughout preference optimization.

090 In detail, our contributions can be summarized as follows:

- 091 • We identify the challenges of dual-model negative preference optimization (NPO), namely its high
092 computational cost and the performance trade-offs inherent in its two separate model design.
- 093 • We demonstrate that simply applying opposing preferences on different branches within a single
094 model leads to poor interactions with DPO-based algorithms, resulting in blurry images as training
095 progresses.
- 096 • We introduce Diff-SNPO, a single-model negative preference optimization method for Diffusion
097 Models. By adapting the BDPO algorithm and deriving an upper bound, our approach effectively
098 eliminates the blurring effect caused by DPO in a negative preference optimization setting.
- 099 • We conduct extensive evaluations on the Pick-a-Pic v2 benchmark with both SD1.5 and SDXL,
100 showing that **Diff-SNPO** delivers strong performance across multiple alignment metrics, while
101 being more efficient in both training and inference.

102 2 RELATED WORK

103 2.1 PREFERENCE OPTIMIZATION

104 Preference optimization aims to align generative models with human judgments. One common strat-
105 egy is to train a *reward model* that scores prompt–image pairs based on semantic or aesthetic quality,

108 and then fine-tune the diffusion model to maximize these scores (Xu et al., 2023; Clark et al., 2024).
 109 Another direction draws on *policy optimization*, which frames denoising as a sequential decision
 110 process: DDPO (Black et al., 2024) applies reinforcement learning across the sampling steps, while
 111 DPOK (Fan et al., 2023) introduces a KL-regularized reward objective. More recently, methods that
 112 bypass explicit reward modeling by learning directly from curated positives or pairwise preferences
 113 have attracted growing interest (Wallace et al., 2024; Lu et al., 2025; Hong et al., 2025). While
 114 these approaches simplify training, they typically handle negative feedback only implicitly, by train-
 115 ing models to favor preferred outputs in relative comparisons. This indirect treatment leaves little
 116 control over explicitly suppressing undesirable generations. In contrast, explicitly modeling dispre-
 117 ferred outcomes offers a more direct form of control: it steers generation away from undesirable
 118 regions of the distribution, reducing artifacts and improving alignment.

119 2.2 NEGATIVE PREFERENCE OPTIMIZATION

120 In response to the limitations of standard preference optimization algorithms, recent work has ex-
 121 plored dual-model strategies that explicitly separate positive and negative preferences. Diff-NPO
 122 (Wang et al., 2025) trains a negative model on inverted preferences and substitutes it for the un-
 123 conditional branch during inference, which improves alignment but introduces a mismatch between
 124 training and inference due to weight interpolation. CHATS (Fu et al., 2025) addresses this issue
 125 by jointly training positive and negative models within a contrastive objective, achieving stronger
 126 alignment but at the cost of increased computation and memory from maintaining two networks.
 127 These challenges highlight the need for approaches that unify positive and negative preference mod-
 128 eling within a single framework. Our method, DIFF-SNPO, takes this direction by integrating both
 129 signals into a shared training objective, reducing overhead while better balancing alignment across
 130 positive and negative preferences.

131 3 PRELIMINARIES

132 3.1 DIFFUSION MODELS

133 Denoising Diffusion Probabilistic Models (DDPMs) (Ho et al., 2020) consist of a forward noising
 134 process and a learned reverse denoising process. The forward process gradually perturbs clean data
 135 x_0 with Gaussian noise according to a variance schedule $\{\beta_t\}_{t=1}^T$:

$$136 \quad q(x_{1:T} | x_0) = \prod_{t=1}^T q(x_t | x_{t-1}), \quad q(x_t | x_{t-1}) = \mathcal{N}\left(x_t; \sqrt{1 - \beta_t} x_{t-1}, \beta_t I\right). \quad (1)$$

137 To invert this process, a neural network ϵ_θ parameterizes the reverse transitions by predicting the
 138 injected noise:

$$139 \quad p_\theta(x_{t-1} | x_t) = \mathcal{N}\left(x_{t-1}; \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(x_t, t) \right), \sigma_t^2 I\right), \quad (2)$$

140 where $\alpha_t = 1 - \beta_t$, $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$, and σ_t^2 denotes the variance.

141 Model training minimizes a variational bound, which simplifies to a weighted noise-prediction loss:

$$142 \quad \mathcal{L}_{\text{DDPM}} = \mathbb{E}_{x_0, \epsilon, t} [\omega(t) \|\epsilon - \epsilon_\theta(x_t, t)\|^2], \quad (3)$$

143 with $\epsilon \sim \mathcal{N}(0, I)$ and $t \sim \mathcal{U}\{1, \dots, T\}$. The weighting function $\omega(t)$ controls the relative con-
 144 tribution of different timesteps, reflecting the varying difficulty of denoising across the diffusion
 145 trajectory. This objective provides a simple and stable training criterion that underpins most modern
 146 diffusion-based generative models.

147 3.2 CLASSIFIER-FREE GUIDANCE

148 Classifier-Free Guidance (CFG) (Ho & Salimans, 2021) modifies the conditional sampling distribu-
 149 tion to strengthen (or weaken) the influence of conditioning information c while remaining anchored
 150 to the unconditional model. At timestep t , the effective distribution is defined as

$$151 \quad \tilde{p}_\theta(x_t | c) \propto p_\theta(x_t) \left(\frac{p_\theta(x_t | c)}{p_\theta(x_t)} \right)^\omega, \quad \omega > 0, \quad (4)$$

162 where $\omega = 1$ recovers the standard conditional model, $\omega > 1$ amplifies the effect of the condition,
 163 and $0 < \omega < 1$ dampens it.

164 In diffusion sampling, the model evolves states using the score function $s_\theta(x_t, t, c) =$
 165 $\nabla_{x_t} \log p_\theta(x_t | c)$. Relating equation 4 to the score yields

$$167 \tilde{s}_\theta(x_t, t, c) = s_\theta(x_t, t) + \omega(s_\theta(x_t, t, c) - s_\theta(x_t, t)), \quad (5)$$

168 where $s_\theta(x_t, t)$ and $s_\theta(x_t, t, c)$ denote the unconditional and conditional scores, respectively. Thus,
 169 CFG moves along the direction that separates the conditional and unconditional scores, with ω
 170 controlling how far we step in that direction.

171 Under the DDPM parameterization Ho et al. (2020) $x_t = \sqrt{\bar{\alpha}_t} x_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon$, the score and the
 172 noise prediction are linked by

$$174 \quad 175 \quad 176 s_\theta(x_t, t, \cdot) = -\frac{1}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(x_t, t, \cdot). \quad (6)$$

177 Substituting equation 6 into equation 5 gives the guided noise estimator

$$178 \quad 179 \tilde{\epsilon}_\theta(x_t, t, c) = \epsilon_\theta(x_t, t) + \omega(\epsilon_\theta(x_t, t, c) - \epsilon_\theta(x_t, t)). \quad (7)$$

180 In essence, CFG modifies the denoiser by amplifying the difference between the conditional and
 181 unconditional predictions, with ω controlling how strongly the generation is steered toward the
 182 conditional distribution. The unconditional prediction $\epsilon_\theta(x_t, t)$ serves as a reference point, and the
 183 correction term $\epsilon_\theta(x_t, t, c) - \epsilon_\theta(x_t, t)$ shifts the generation trajectory away from regions favored
 184 by the unconditional model but inconsistent with the conditioning signal c . This allows CFG to
 185 suppress undesirable or generic outputs that the unconditional model might produce on its own, en-
 186 couraging samples that better reflect the intended conditioning. As a result, CFG offers a principled
 187 and tunable mechanism to improve text-image alignment.

188 3.3 DIFFUSION DIRECT PREFERENCE OPTIMIZATION

190 Preference alignment is often formalized using the Bradley–Terry (BT) model (Kendall & Smith,
 191 1940). Building on this idea, Direct Preference Optimization (DPO) (Rafailov et al., 2023) sidesteps
 192 the need to fit an explicit reward model by defining an implicit reward through likelihood ratios
 193 between the policy π_θ and a reference policy π_{ref} :

$$194 \quad 195 \quad 196 \mathcal{L}_{\text{DPO}}(\theta) = -\mathbb{E}_{(x_0^w, x_0^l, c) \sim \mathcal{D}} \left[\log \sigma \left(\beta \left(\log \frac{\pi_\theta(x_0^w | c)}{\pi_{\text{ref}}(x_0^w | c)} - \log \frac{\pi_\theta(x_0^l | c)}{\pi_{\text{ref}}(x_0^l | c)} \right) \right) \right]. \quad (8)$$

197 Here, $\beta > 0$ is a temperature parameter that controls the strength of the preference signal, and
 198 $\sigma(\cdot)$ denotes the sigmoid function. Extending this objective to diffusion models requires assigning
 199 preferences over entire trajectories $x_{0:T}$. In this context, Diff-DPO (Wallace et al., 2024) derives an
 200 upper bound on the exact Diff-DPO objective:

$$202 \quad 203 \quad 204 \mathcal{L}_{\text{Diff-DPO}}(\theta; y) = -\mathbb{E}_{(x_0^w, x_0^l, c) \sim \mathcal{D}, t \sim \mathcal{U}[1, T]} \left[\log \sigma \left(y T \omega(t) \beta (\Delta_t^w(c) - \Delta_t^l(c)) \right) \right], \quad (9)$$

205 with

$$206 \quad 207 \quad 208 \Delta_t^w(c) = \left\| \epsilon^w - \epsilon_\theta(x_t^w, t, c) \right\|_2^2 - \left\| \epsilon^w - \epsilon_{\text{ref}}(x_t^w, t, c) \right\|_2^2, \quad (10)$$

$$209 \quad 210 \quad 211 \Delta_t^l(c) = \left\| \epsilon^l - \epsilon_\theta(x_t^l, t, c) \right\|_2^2 - \left\| \epsilon^l - \epsilon_{\text{ref}}(x_t^l, t, c) \right\|_2^2, \quad (11)$$

212 Here, x_t^w and x_t^l are noisy states obtained by applying the forward diffusion process to the clean
 213 samples x_0^w, x_0^l with corresponding noise terms $\epsilon^w, \epsilon^l \sim \mathcal{N}(0, I)$. The label $y \in \{+1, -1\}$ encodes
 214 the preference direction and $\omega(t)$ is a time-dependent weighting function.

215 3.4 DIFFUSION NEGATIVE PREFERENCE OPTIMIZATION

216 Diffusion Negative Preference Optimization (Diff-NPO) (Wang et al., 2025) extends standard pref-
 217 erence learning by explicitly modeling undesirable behavior. It does so by training on an *inverted*

preference dataset, where the roles of winners and losers are swapped. This yields a negatively aligned model, θ^- , which learns to assign higher likelihood to dispreferred samples. During sampling, Diff-NPO replaces the standard unconditional branch in classifier-free guidance with this negative model, leading to the following guidance formulation:

$$\epsilon_{\text{NPO}}(x_t, c) = \epsilon_{\theta^+}(x_t, c) + \omega [\epsilon_{\theta^+}(x_t, c) - \epsilon_{\theta^-}(x_t, \emptyset)], \quad (12)$$

where ϵ_{θ^+} is the positively aligned model trained on standard preference data, and ω is a guidance strength hyperparameter.

In practice, however, the positive and negative models that are trained independently often exhibit poor correlation, which undermines the effectiveness of ϵ_{θ^-} as a meaningful contrastive signal during sampling. To mitigate this, Diff-NPO applies a weight-merging procedure that combines the reference model θ_{ref} , the positive model θ^+ , and the negative model θ^- :

$$\hat{\theta}^- = \theta_{\text{ref}} + \alpha(\theta^+ - \theta_{\text{ref}}) + \beta(\theta^- - \theta_{\text{ref}}), \quad (13)$$

where α and β are interpolation coefficients. In practice, the merged parameters $\hat{\theta}^-$ replace ϵ_{θ^-} in the guidance formulation of Eq. 13, yielding improved generation quality. However, this benefit comes at the cost of higher training overhead and weakened negative alignment.

4 METHOD

4.1 LIMITATIONS OF NEGATIVE PREFERENCE OPTIMIZATION

Dual-model approaches to Negative Preference Optimization (NPO) (Wang et al., 2025) offer direct way to model positive and negative preferences, training separate models— θ^+ and θ^- —for each. However, this design introduces significant practical overhead. Maintaining two networks doubles memory and compute requirements, increases training time. These costs quickly escalate with large diffusion backbones, making dual-model setups difficult to scale in practice.

At inference time, using the two models separately preserves the intended contrast between positive and negative preferences. However, this approach often compromises sample quality, especially when generating from the negatively aligned model θ^- . To mitigate this, Diff-NPO introduces a merged parameter formulation, $\hat{\theta}^-$, which interpolates between θ^+ , θ^- , and the reference backbone θ_{ref} as shown in Eq. 13. This interpolation enhances generation quality and increases agreement between the two models’ predictions (Wang et al., 2025), but it also biases the merged model toward θ^+ , reducing the influence of θ^- . As illustrated in Figure 1, this trade-off between fidelity and alignment exposes a deeper limitation of the dual-model paradigm: its reliance on separate networks and post-hoc merging undermines the very preference contrast that NPO is designed to enforce. To overcome these issues, we explore single-model formulations that unify both positive and negative preference signals within a shared architecture—offering a more robust and scalable solution, which we introduce in the next section.

4.2 ISSUES WITH NAIVE DIFF-SNPO

Standard diffusion models inherently support both conditional and unconditional branches within a single network—a design popularized by Classifier-Free Guidance (Ho & Salimans, 2021). This built-in dual-branch structure presents a natural alternative to the dual-model Diff-NPO setup: rather than maintaining two separate networks, one can leverage the existing architecture by updating the conditional branch with preferred samples and the unconditional (or negatively conditioned) branch with dispreferred ones. This approach preserves the contrast between preference signals while avoiding the redundancy and overhead of training and managing separate models.

Specifically, let $Y \in \{+1, -1\}$ be a branch label with $\Pr(Y = +1) = 1 - p$ and $\Pr(Y = -1) = p$, where

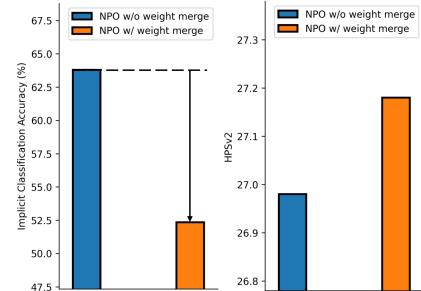


Figure 1: **Negative implicit accuracy (left) and HPSv2 (right) on SD1.5.** Weight merging lowers implicit accuracy while increasing reward, revealing a trade-off.

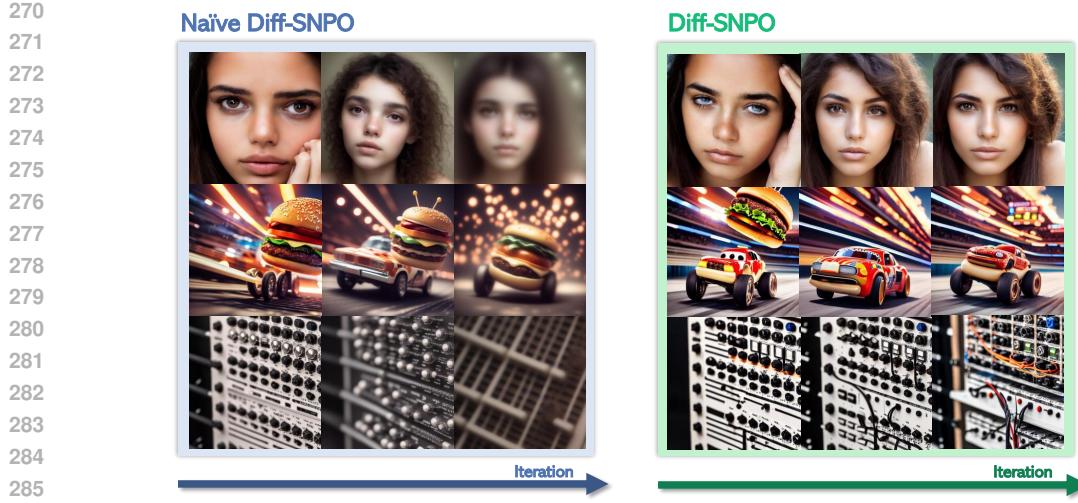


Figure 3: **Generated samples across training iterations.** Naïve-SNPO produces progressively blurred outputs during training, while Diff-SNPO remains stable and yields increasingly preference-aligned images without blurring artifacts.

$p \in [0, 1]$ is the CFG dropout probability. The effective conditioning is

$$\tilde{c}(Y) = \begin{cases} c, & Y = +1 \text{ (conditional branch)} \\ \emptyset, & Y = -1 \text{ (unconditional/null branch)}, \end{cases}$$

where c denotes the conditioning input (e.g., a text prompt) and \emptyset indicates null conditioning, as in CFG. A Naive Diff-SNPO objective can thus be constructed by adapting the Diff-DPO objective in Eq. 9:

$$\mathcal{L}_{\text{Naive Diff-SNPO}}(\theta) = -\mathbb{E}_{(x_0^w, x_0^l) \sim \mathcal{D}, t, Y} \left[\log \sigma(Y T \omega(t) \beta (\Delta_t^w(\tilde{c}(Y)) - \Delta_t^l(\tilde{c}(Y)))) \right], \quad (14)$$

While Naive Diff-SNPO provides a simple and intuitive way to incorporate flipped preferences within a single network, it exhibits substantial degradation in generation quality over training. As shown in Fig. 3, this naive objective produces increasingly blurry outputs, with reduced contrast and diminished high-frequency detail. We attribute this behavior to a known property of the DPO objective: improvements in pairwise margin often come from decreasing the likelihood of both candidates—penalizing the loser more strongly—rather than consistently increasing the likelihood of the winner (Pal et al., 2024; Rafailov et al., 2024b; Cho et al., 2025). This trend is reflected in our results as well, as illustrated in Fig. 2.

To quantify this effect on preferred samples, we track the win-sample likelihood ratio against a fixed reference model, given by:

$$\frac{\pi_\theta(x^w)}{\pi_{\text{ref}}(x^w)} \approx \mathbb{E}_{x_0^w, t} \left[e^{\Delta_t^w(c)} \right].$$

Under Naive Diff-SNPO, we observe a consistent decrease in the relative win probability of preferred samples over the course of training (Fig. 3), indicating that the model suppresses likelihoods across both branches rather than reinforcing the preferred ones. This undesirable trend, combined with the structure of the optimization signals, leads to a characteristic blurring effect. Specifically, Naive Diff-SNPO applies symmetric but opposing updates to the conditional and unconditional branches—encouraging one to increase the likelihood of a sample while the other decreases it. Since both branches share model parameters, these conflicting gradients interfere with each other, and the model resolves this

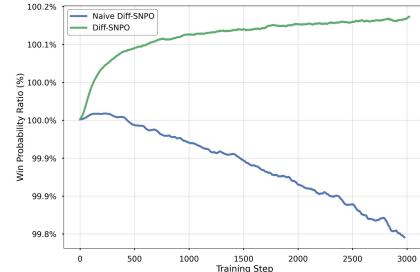


Figure 2: **Win probability ratio over training.** Naïve-SNPO’s win probability decreases during training, whereas Diff-SNPO’s improves steadily.

324 conflict by averaging them. In generative settings, this averaging dampens contrast and degrades
 325 fine detail, resulting in blurry outputs. To overcome this limitation, we seek to break the destructive
 326 symmetry in the update rule. In the next section, we introduce an asymmetric preference optimiza-
 327 tion approach that biases learning toward increasing win probability, thereby mitigating gradient
 328 interference and addressing the blurring effect.

330 4.3 DIFF-SNPO

331 Recently, Cho et al. (2025) introduced Bounded DPO (BDPO) to address a key shortcoming of
 332 standard DPO. As the model reduces the probability assigned to the losing sample, the corresponding
 333 lose sample log-likelihood term, $\log \pi_\theta(\mathbf{y}_l \mid \mathbf{x})$, in the objective grows disproportionately large,
 334 causing the loss to become dominated by the loser. This skews the gradient signal and can even
 335 drive updates that decrease the likelihood of the preferred (winning) output, despite the preference
 336 label.

337 To mitigate this issue, BDPO replaces the losing term with a mixture distribution that includes a
 338 non-vanishing contribution from the reference policy:

$$339 \pi_{\text{mix}}(\mathbf{y} \mid \mathbf{x}) = \lambda \pi_\theta(\mathbf{y} \mid \mathbf{x}) + (1 - \lambda) \pi_{\text{ref}}(\mathbf{y} \mid \mathbf{x}), \quad \lambda \in (0, 1),$$

340 which leads to the modified objective

$$341 \mathcal{L}_{\text{BDPO}}(\pi_\theta; \pi_{\text{ref}}) = -\mathbb{E}_{(x_0^w, x_0^l, c) \sim \mathcal{D}} \left[\log \sigma \left(\beta \left[\log \frac{\pi_\theta(x_0^w \mid c)}{\pi_{\text{ref}}(x_0^w \mid c)} - \log \frac{\pi_{\text{mix}}(x_0^l \mid c)}{\pi_{\text{ref}}(x_0^l \mid c)} \right] \right) \right]. \quad (15)$$

342 This modification bounds the contribution of the loser term and prevents it from overwhelming
 343 the loss, thereby preserving the intended effect of preference optimization: promoting the winning
 344 sample. In addition, BDPO retains the same global minimizers as DPO while enforcing a lower
 345 bound on the winning likelihood, offering stronger stability guarantees throughout training (Cho
 346 et al., 2025).

347 To adapt BDPO to diffusion models, similar to Diff-DPO, we define the trajectory-level reward

$$348 r_\theta(c, \mathbf{x}_0) = \mathbb{E}_{\mathbf{x}_{1:T} \sim p_\theta(\cdot \mid \mathbf{x}_0, c)} \left[R(c, \mathbf{x}_{0:T}) \right]. \quad (16)$$

349 The corresponding Diffusion BDPO objective is

$$350 \mathcal{L}_{\text{Diff-BDPO}}(\theta) = -\mathbb{E}_{(\mathbf{x}_{0:T}^w, \mathbf{x}_{0:T}^l, c) \sim \mathcal{D}} \left[\log \sigma \left(\beta \left[\log \frac{p_\theta(\mathbf{x}_{0:T}^w \mid c)}{p_{\text{ref}}(\mathbf{x}_{0:T}^w \mid c)} - \log \frac{\pi_{\text{mix}}(\mathbf{x}_{0:T}^l \mid c)}{p_{\text{ref}}(\mathbf{x}_{0:T}^l \mid c)} \right] \right) \right]. \quad (17)$$

351 Following Diff-DPO, we upper-bound this objective via its ELBO and then apply Jensen's inequality
 352 to obtain the result below (see Appendix A.2 for the full derivation):

$$353 \mathcal{L}_{\text{Diff-BDPO-UB}}(\theta) = -\mathbb{E}_{(x_0^w, x_0^l, c) \sim \mathcal{D}, t \sim \mathcal{U}[1, T]} \left[\log \sigma \left(-\beta (m(x_t^w, c) - m_{\text{mix}}(x_t^l, c)) \right) \right]. \quad (18)$$

354 Here the per-step terms are:

$$355 d_\theta(x_t, \epsilon, t, c) = T \omega(t) \|\epsilon - \epsilon_\theta(x_t, t, c)\|_2^2, \quad d_{\text{ref}}(x_t, \epsilon, t, c) = T \omega(t) \|\epsilon - \epsilon_{\text{ref}}(x_t, t, c)\|_2^2, \quad (19)$$

$$356 m(x_t, c) = d_\theta(x_t, \epsilon, t, c) - d_{\text{ref}}(x_t, \epsilon, t, c), \quad (20)$$

$$357 m_{\text{mix}}(x_t, c) = -\log \left(\lambda e^{-d_\theta(x_t, \epsilon, t, c)} + (1 - \lambda) e^{-d_{\text{ref}}(x_t, \epsilon, t, c)} \right) - d_{\text{ref}}(x_t, \epsilon, t, c). \quad (21)$$

358 where $\mathcal{L}_{\text{Diff-BDPO-UB}}$ denotes the upper bound approximation of Diff-BDPO.

359 Building on this, we define our final Diff-SNPO objective, which applies our introduced Diff-BDPO-
 360 UB objective into a single model negative preference optimization framework:

$$361 \mathcal{L}_{\text{SNPO}}(\theta) = -\mathbb{E}_{(\mathbf{x}^w, \mathbf{x}^l, c) \sim \mathcal{D}, t \sim p(t), Y} \left[\log \sigma \left(\beta (m(\tilde{x}^w(Y), \tilde{c}(Y)) - m_{\text{mix}}(\tilde{x}^l(Y), \tilde{c}(Y))) \right) \right]. \quad (22)$$

362 where $\tilde{x}^{w/l}(Y)$ is given as:

$$363 \tilde{x}^{w/l}(Y) = \begin{cases} x^{w/l}, & \text{if } Y = +1 \quad (\text{conditional branch}) \\ x^{l/w}, & \text{if } Y = -1 \quad (\text{unconditional/null branch}). \end{cases} \quad (23)$$

364 As shown in Fig. 3, Diff-SNPO avoids the decline in winner likelihood observed with Naive Diff-
 365 SNPO: the estimated log-likelihood of winning samples steadily improves throughout training. In
 366 line with this, Fig. 3 further shows that Diff-SNPO prevents the progressive low-contrast “blurring”
 367 artifact characteristic of the naive training objective, effectively addressing its key shortcoming.

378
 379 Table 2: **Comparison of Diff-SNPO with baseline methods on SD1.5 and SDXL backbones**
 380 **using Pick-a-Pic v2.** All values are reported as mean \pm 95% confidence interval over 4 random
 381 seeds. Diff-SNPO consistently achieves the highest scores across most human preference metrics,
 382 reflecting improved alignment and visual quality. For clarity, the best-performing method in each
 383 metric is shown in **bold**, and the second-best is underlined.

Model	Method	HPSv2	Pick Score	Aesthetic Score	Image Reward
SD1.5	Baseline	26.24 ± 0.14	20.64 ± 0.09	5.2849 ± 0.12	0.1221 ± 0.08
	Diff-DPO (Wallace et al., 2024)	26.55 ± 0.04	21.01 ± 0.04	5.3823 ± 0.01	0.2968 ± 0.07
	Diff-NPO (Wang et al., 2025)	26.92 ± 0.05	21.46 ± 0.04	5.5381 ± 0.04	0.3786 ± 0.09
	CHATS (Fu et al., 2025)	27.20 ± 0.16	21.05 ± 0.06	5.6845 ± 0.07	0.2995 ± 0.07
	Diff-BDPO (Ours)	26.64 ± 0.16	21.15 ± 0.10	5.4457 ± 0.10	0.3165 ± 0.11
	Diff-SNPO (Ours)	27.23 ± 0.07	22.24 ± 0.03	5.6258 ± 0.04	0.6936 ± 0.07
SDXL	Baseline	27.43 ± 0.07	22.13 ± 0.06	5.8850 ± 0.01	0.7605 ± 0.05
	Diff-DPO (Wallace et al., 2024)	28.09 ± 0.03	22.59 ± 0.02	5.8884 ± 0.01	0.9841 ± 0.07
	Diff-NPO (Wang et al., 2025)	28.30 ± 0.08	22.67 ± 0.05	5.9449 ± 0.02	0.9847 ± 0.01
	CHATS (Fu et al., 2025)	28.25 ± 0.10	22.34 ± 0.08	5.8785 ± 0.03	1.0543 ± 0.03
	Diff-BDPO (Ours)	28.13 ± 0.04	22.36 ± 0.10	5.8037 ± 0.07	0.9946 ± 0.05
	Diff-SNPO (Ours)	28.33 ± 0.08	22.69 ± 0.06	5.8129 ± 0.06	1.0100 ± 0.06

396 5 EXPERIMENTS

397 5.1 EXPERIMENTAL SETUP

400 **Datasets and Models.** We fine-tune both Stable Diffusion v1.5 (SD1.5) (Rombach et al., 2022)
 401 and Stable Diffusion XL (SDXL) (Podell et al., 2024) using our proposed Diffusion SNPO objective
 402 in Eq. 22. For training, we utilize the Pick-a-Pic v2 (Kirstain et al., 2023) corpus, a large-scale
 403 human preference dataset comprising 851,293 image pairs across 58,960 unique prompts.

405 **Training Details.** Models are initialized from the publicly available SD1.5 (Rombach et al., 2022)
 406 (CreativeML Open RAIL-M license) and SDXL (Podell et al., 2024) (MIT license) checkpoints.
 407 Training is conducted using the AdamW (Loshchilov & Hutter, 2019) optimizer with a learning rate
 408 of 2.048×10^{-8} . We use a batch size of 512 and train for 3,000 steps on SD1.5, and a larger batch
 409 size of 2,048 for 625 steps on SDXL. All experiments are run on $8 \times$ NVIDIA A6000 GPUs using
 410 distributed data parallelism. Total training time is approximately 12 hours for SD1.5 and 68 hours
 411 for SDXL. The regularization coefficient is set to $\beta = 2000$ for SD1.5 and $\beta = 5000$ for SDXL.
 412 Unless stated otherwise, Diff-SNPO is trained with an interpolation parameter of $\lambda = 0.9$.

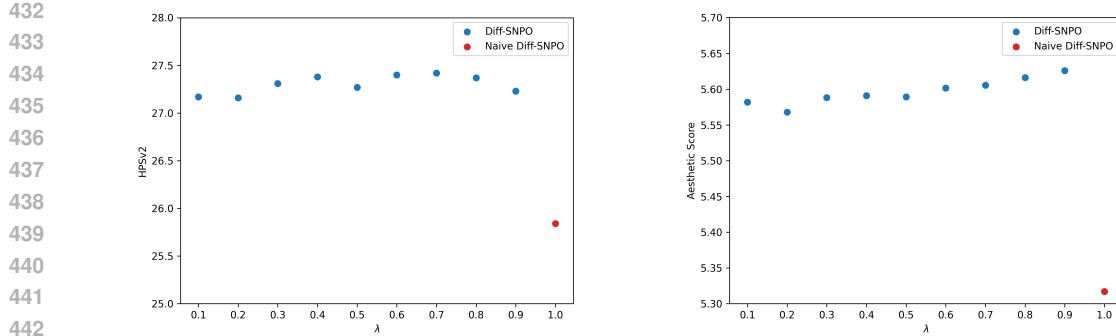
413 **Baselines.** We compare Diff-SNPO against both negative preference optimization methods and
 414 standard alignment baselines. Specifically, we evaluate Diff-NPO (Wang et al., 2025), CHATS (Fu
 415 et al., 2025), and Diff-DPO (Wallace et al., 2024), along with the original pretrained models for
 416 SD1.5 (Rombach et al., 2022) and SDXL (Podell et al., 2024).

418 **Evaluation Protocol.** All models are evaluated using the DDIM (Song et al., 2021) sampler with
 419 50 inference steps and a classifier-free guidance scale of 7.5. For Diff-NPO and CHATS, we adopt
 420 the hyperparameters specified in their released code, as these methods modify the sampling pipeline.
 421 Performance is assessed using five widely adopted reward models: PickScore (Kirstain et al., 2023),
 422 HPSv2 (Wu et al., 2023), ImageReward (Xu et al., 2023), and Aesthetics Score (Schuhmann, 2023).
 423 To account for stochasticity, we report average results across four random seeds, with each model
 424 generating 2,000 images using prompts from the Pick-a-Pic v2 test set.

425 5.2 QUANTITATIVE RESULTS

427 Table 2 presents quantitative comparisons of Diff-SNPO against established baselines on SD1.5 and
 428 SDXL backbones.

429 On SD1.5, Diff-SNPO consistently outperforms prior methods across most human preference met-
 430 rics. Notably, it achieves 27.23 on HPSv2 and 22.24 on PickScore, marking a substantial improve-
 431 ment over both Diff-DPO and the more stable Diff-BDPO. This indicates that the gains stem not just

Figure 4: **Ablation over λ on SD 1.5 for both HPSv2 and Aesthetic score.**

from BDPO’s stabilization, but from negative preference optimization itself. While CHATS attains a slightly higher Aesthetic Score, it lags behind on all other metrics, suggesting that its more visually pleasing outputs come at the cost of semantic alignment. Diff-SNPO also surpasses its dual-model counterpart, Diff-NPO, across all metrics while using only half the computation. We attribute this to its single-model design, which preserves negative preference alignment better than Diff-NPO, where accuracy drops notably after model merging.

On SDXL, Diff-SNPO delivers strong and competitive results, with HPSv2 and ImageReward scores of 28.33 and 1.01, comparable to state-of-the-art methods. Its Aesthetic Score (5.81) is slightly below DPO (5.89), and the large advantage it shows over Diff-NPO on SD1.5 becomes smaller on SDXL. This stems from Diff-SNPO’s bias toward optimizing “preferred” samples in the dataset. Although these samples are favored over their losing counterparts, they have been shown to be less aesthetically pleasing than the already strong outputs of the base SDXL model. As a result, focusing on them may slightly reduce aesthetic quality and limit gains on stronger backbones. Similar trends appear with CHATS and Diff-BDPO, suggesting the limitation comes from the dataset rather than the alignment method itself. Even so, Diff-SNPO retains important practical benefits: it requires only half the computation of Diff-NPO and, with its single-model architecture, also enables faster sampling—making it both efficient and scalable.

5.3 IMPACT OF THE MIXING PARAMETERS λ AND β ON DIFF-SNPO

In this section, we examine the impact of the parameters λ and β on Diff-SNPO using SD 1.5, while keeping all other hyperparameters identical to those in the main experiments.

From Fig. 4, we observe that Diff-SNPO remains stable as λ varies, showing only minor fluctuations across both reward metrics. This indicates that, within the tested range, the choice of λ has limited impact on performance. In contrast, the setting $\lambda = 1.0$, corresponding to Naive-SNPO, shows a pronounced drop in both reward metrics. As discussed in Section 4.2, this decline stems from the blurring artifacts that emerge when the win likelihood decreases under single-model Negative Preference Optimization, ultimately degrading image quality.

When varying β , Table 3 shows that performance is similarly stable across different choices. Therefore, we retain the β values from the original Diff-DPO configuration, as they provide a reliable default without introducing meaningful variability in performance.

5.4 COMPARING NEGATIVE PREFERENCE ALIGNMENT ACROSS METHODS

From Table 5, we observe that Diff-DPO and Diff-BDPO reach similar negative implicit accuracy in their unconditional branches, both consistently below 50%. This behavior is expected because their unconditional branches are trained only with positive preference alignment, which limits their ability to learn negative preference signals. In contrast, Diff-NPO experiences a substantial drop in negative implicit accuracy after weight merging, decreasing by more than 10%. Its post-merge accuracy also falls below that of Diff-SNPO and lies several points behind its own pre-merged value.

β	HPSv2	Aesthetic Score
1000	27.30	5.6682
2000	27.23	5.6258
3000	27.23	5.5886

Table 3: Effect of β on the Pick-a-Pic v2 dataset.

486
 487 Table 4: **Training cost comparison.** Experiments were conducted on $8 \times$ A6000 GPUs with a total
 488 batch size of 512. Dual-model approaches require substantially more memory and incur slower
 489 training throughput. The best result in each column is shown in **bold**.

Method	Memory (GB) \downarrow	Time / Step (s) \downarrow	Relative Speed \uparrow
Diff-NPO	44.2×2	12.25×2	$1.00 \times$
CHATS	46.3	13.19	$1.86 \times$
Diff-SNPO (Ours)	44.2	12.25	2.00 \times

490
 491 Table 5: **Negative preference implicit classification accuracy and loss.** Parentheses denote Diff-
 492 NPO without weight merging. Diff-SNPO achieves higher negative implicit accuracy and lower
 493 negative preference loss than Diff-NPO, improving its negative alignment.

Method	Neg. Implicit Acc. (%)	Neg. Pref. Loss
Diff-DPO	31.86	0.759
Diff-BDPO	32.04	0.805
Diff-NPO	52.34 (63.80)	0.703 (0.648)
Diff-SNPO	57.45	0.668

504
 505 Although implicit accuracy is not a definitive measure of overall performance, and can be influenced
 506 by reward-model artifacts (Rafailov et al., 2024a; Amini et al., 2024), it is still useful for revealing
 507 the sampling and training mismatch present in Diff-NPO. This mismatch is further reflected in the
 508 negative preference alignment loss, which rises sharply after weight merging. Diff-SNPO does not
 509 suffer from this problem: its single-model design preserves the negative alignment signal and avoids
 510 the degradation observed in Diff-NPO. As a result, Diff-SNPO maintains more reliable negative
 511 preference modeling and more effectively steers generation away from undesirable samples.

5.5 TRAINING COMPUTATIONAL COST

522 Beyond alignment performance, the practicality of a preference optimization algorithm also depends
 523 on its computational efficiency. To assess this, we compare the training computational cost of dual-
 524 model approaches (CHATS and Diff-NPO) against our single-model Diff-SNPO. As reported in
 525 Table 4, the single-model design of Diff-SNPO yields substantial efficiency gains in both memory
 526 usage and training time. By eliminating the need to train two separate networks, Diff-SNPO reduces
 527 memory consumption and achieves a $2 \times$ speedup in per-step training time relative to its dual-model
 528 counterpart, Diff-NPO. A similar advantage is observed over CHATS: while both methods require
 529 comparable memory, Diff-SNPO trains faster because it optimizes the conditional and unconditional
 530 branches in parallel, whereas CHATS processes them sequentially. In summary, Diff-SNPO com-
 531 bines lower memory overhead with faster training, establishing it as a more efficient and scalable
 532 alternative for negative preference optimization. *A detailed comparison of inference cost, which*
 533 *further highlights the efficiency of our single-model approach, can be found in Appendix A.7.*

6 CONCLUSION

534 In conclusion, we propose Diff-SNPO, a single-model framework for Negative Preference Opti-
 535 mization that achieves strong performance while simplifying the training pipeline. In contrast to
 536 prior approaches that require two separate models, Diff-SNPO unifies conditional and unconditional
 537 branches within a single architecture, thereby eliminating redundant computation and substantially
 538 improving efficiency in both training and inference. Our experiments demonstrate that this stream-
 539 lined design not only matches, but often exceeds the performance of existing methods, highlighting
 540 its effectiveness in preserving negative preference alignment. Beyond raw performance, the reduced
 541 computational footprint and faster sampling make Diff-SNPO a practical and scalable solution, low-
 542 ering the barriers to applying preference optimization in real-world generative modeling tasks.

540 REFERENCES
541

542 Afra Amini, Tim Vieira, and Ryan Cotterell. Direct preference optimization with an offset. In
543 Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Findings of the Association for Com-
544 putational Linguistics: ACL 2024*, pp. 9954–9972, Bangkok, Thailand, August 2024. Associa-
545 tion for Computational Linguistics. doi: 10.18653/v1/2024.findings-acl.592. URL <https://aclanthology.org/2024.findings-acl.592/>.

546

547 Kevin Black, Michael Janner, Yilun Du, Ilya Kostrikov, and Sergey Levine. Training diffusion
548 models with reinforcement learning. In *The Twelfth International Conference on Learning Rep-
549 resentations*, 2024. URL <https://openreview.net/forum?id=YCWjhGrJFD>.

550

551 Jay Hyeon Cho, JunHyeok Oh, Myunsoo Kim, and Byung-Jun Lee. Rethinking dpo: The role
552 of rejected responses in preference misalignment, 2025. URL <https://arxiv.org/abs/2506.12725>.

553

554 Kevin Clark, Paul Vicol, Kevin Swersky, and David J. Fleet. Directly fine-tuning diffusion models
555 on differentiable rewards. In *The Twelfth International Conference on Learning Representations*,
556 2024. URL <https://openreview.net/forum?id=1vmSEVL19f>.

557

558 Ying Fan, Olivia Watkins, Yuqing Du, Hao Liu, Moonkyung Ryu, Craig Boutilier, Pieter Abbeel,
559 Mohammad Ghavamzadeh, Kangwook Lee, and Kimin Lee. Reinforcement learning for fine-
560 tuning text-to-image diffusion models. In *Thirty-seventh Conference on Neural Information Pro-
561 cessing Systems*, 2023. URL <https://openreview.net/forum?id=80TPepXzeh>.

562

563 Minghao Fu, Guo-Hua Wang, Liangfu Cao, Qing-Guo Chen, Zhao Xu, Weihua Luo, and Kaifu
564 Zhang. CHATS: Combining human-aligned optimization and test-time sampling for text-to-image
565 generation. In *Forty-second International Conference on Machine Learning*, 2025. URL <https://openreview.net/forum?id=D4Y71nbGRg>.

566

567 Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on
568 Deep Generative Models and Downstream Applications*, 2021. URL <https://openreview.net/forum?id=qw8AKxfYbI>.

569

570 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In
571 H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neu-
572 ral Information Processing Systems*, volume 33, pp. 6840–6851. Curran Associates, Inc.,
573 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf.

574

575 Jonathan Ho, Tim Salimans, Alexey A. Gritsenko, William Chan, Mohammad Norouzi, and
576 David J. Fleet. Video diffusion models. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave,
577 and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022. URL
578 https://openreview.net/forum?id=f3zNgKga_ep.

579

580 Jiwoo Hong, Sayak Paul, Noah Lee, Kashif Rasul, James Thorne, and Jongheon Jeong. Margin-
581 aware preference optimization for aligning diffusion models without reference. In *First Workshop
582 on Scalable Optimization for Efficient and Adaptive Foundation Models*, 2025. URL <https://openreview.net/forum?id=JMVWaMR53s>.

583

584 Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of
585 diffusion-based generative models. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and
586 Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022. URL
587 <https://openreview.net/forum?id=k7FuTOWMOC7>.

588

589 M. G. Kendall and B. Babington Smith. On the method of paired comparisons. *Biometrika*, 31
590 (3-4):324–345, 03 1940. ISSN 0006-3444. doi: 10.1093/biomet/31.3-4.324. URL <https://doi.org/10.1093/biomet/31.3-4.324>.

591

592 Yuval Kirstain, Adam Polyak, Uriel Singer, Shahbuland Matiana, Joe Penna, and Omer Levy. Pick-
593 a-pic: An open dataset of user preferences for text-to-image generation. 2023.

594 Shufan Li, Konstantinos Kallidromitis, Akash Gokul, Yusuke Kato, and Kazuki Kozuka. Aligning
 595 diffusion models by optimizing human utility. In *The Thirty-eighth Annual Conference on Neu-*
 596 *ral Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=MTMShU5QaC>.

598 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Confer-*
 599 *ence on Learning Representations*, 2019. URL <https://openreview.net/forum?id=Bkg6RiCqY7>.

602 Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. DPM-solver: A
 603 fast ODE solver for diffusion probabilistic model sampling in around 10 steps. In Alice H. Oh,
 604 Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information*
 605 *Processing Systems*, 2022. URL https://openreview.net/forum?id=2uAaGwlP_V.

606 Yunhong Lu, Qichao Wang, Hengyuan Cao, Xierui Wang, Xiaoyin Xu, and Min Zhang. Inpo: Inver-
 607 sion preference optimization with reparametrized ddim for efficient diffusion model alignment. In
 608 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*,
 609 pp. 28629–28639, June 2025.

610 Arka Pal, Deep Karkhanis, Samuel Dooley, Manley Roberts, Siddartha Naidu, and Colin White.
 611 Smaug: Fixing failure modes of preference optimisation with dpo-positive. *arXiv preprint*
 612 *arXiv:2402.13228*, 2024.

614 Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe
 615 Penna, and Robin Rombach. SDXL: Improving latent diffusion models for high-resolution image
 616 synthesis. In *The Twelfth International Conference on Learning Representations*, 2024. URL
 617 <https://openreview.net/forum?id=di52zR8xgf>.

618 Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea
 619 Finn. Direct preference optimization: Your language model is secretly a reward model. In
 620 *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=HPuSIXJaa9>.

622 Rafael Rafailov, Yaswanth Chittepu, Ryan Park, Harshit Sikchi, Joey Hejna, Bradley Knox, Chelsea
 623 Finn, and Scott Niekum. Scaling laws for reward model overoptimization in direct alignment
 624 algorithms, 2024a. URL <https://arxiv.org/abs/2406.02900>.

626 Rafael Rafailov, Joey Hejna, Ryan Park, and Chelsea Finn. From \$r\$ to \$q^*\$: Your language
 627 model is secretly a q-function. In *First Conference on Language Modeling*, 2024b. URL <https://openreview.net/forum?id=kEVcNxtqXk>.

629 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
 630 resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Con-*
 631 *ference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10684–10695, June 2022.

632 Ludan Ruan, Yiyang Ma, Huan Yang, Huiguo He, Bei Liu, Jianlong Fu, Nicholas Jing Yuan, Qin
 633 Jin, and Baining Guo. Mm-diffusion: Learning multi-modal diffusion models for joint audio and
 634 video generation. In *CVPR*, 2023.

636 Patrick Schramowski, Manuel Brack, Björn Deiseroth, and Kristian Kersting. Safe latent diffusion:
 637 Mitigating inappropriate degeneration in diffusion models. In *Proceedings of the IEEE Confer-*
 638 *ence on Computer Vision and Pattern Recognition (CVPR)*, 2023.

639 Christoph Schuhmann. Laion-aesthetics. <https://github.com/christophschuhmann/improved-aesthetic-predictor>, October 2023.

641 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *Inter-*
 642 *national Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=St1giarCHLP>.

645 Bram Wallace, Meihua Dang, Rafael Rafailov, Linqi Zhou, Aaron Lou, Senthil Purushwalkam,
 646 Stefano Ermon, Caiming Xiong, Shafiq Joty, and Nikhil Naik. Diffusion model alignment using
 647 direct preference optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision*
 and *Pattern Recognition (CVPR)*, pp. 8228–8238, June 2024.

648 Fu-Yun Wang, Yunhao Shui, Jingtian Piao, Keqiang Sun, and Hongsheng Li. Diffusion-NPO:
649 Negative preference optimization for better preference aligned generation of diffusion mod-
650 els. In *The Thirteenth International Conference on Learning Representations*, 2025. URL
651 <https://openreview.net/forum?id=iJi7nz5Cxc>.

652 Xiaoshi Wu, Yiming Hao, Keqiang Sun, Yixiong Chen, Feng Zhu, Rui Zhao, and Hongsheng Li.
653 Human preference score v2: A solid benchmark for evaluating human preferences of text-to-
654 image synthesis. *arXiv preprint arXiv:2306.09341*, 2023.

655 Xun Wu, Shaohan Huang, Lingjie Jiang, and Furu Wei. Rethinking dpo-style diffusion aligning
656 frameworks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*
(ICCV), pp. 18068–18077, October 2025.

657 Jiazheng Xu, Xiao Liu, Yuchen Wu, Yuxuan Tong, Qinkai Li, Ming Ding, Jie Tang, and Yuxiao
658 Dong. Imagereward: Learning and evaluating human preferences for text-to-image generation.
659 In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=JVzeOYEx6d>.

660 Wenliang Zhao, Lujia Bai, Yongming Rao, Jie Zhou, and Jiwen Lu. UniPC: A unified predictor-
661 corrector framework for fast sampling of diffusion models. In *Thirty-seventh Conference on*
662 *Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=hrkmlPhp1u>.

663

664

665

666

667

668

669

670

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702 **A APPENDIX**
 703

704 **A.1 THE USE OF LLMs**
 705

706 We used LLMs solely for light editing such as correcting grammatical errors and polishing some
 707 words. They did not contribute to research ideation, experiments, analysis, or substantive writing.
 708

709 **A.2 DERIVATION OF THE DIFFUSION-BDPO UPPER BOUND**
 710

711 We begin by recalling the Diffusion-BDPO objective:

712
$$\mathcal{L}_{\text{Diff-BDPO}}(\theta) = -\mathbb{E}_{(\mathbf{x}_{0:T}^w, \mathbf{x}_{0:T}^l, c) \sim \mathcal{D}} \left[\log \sigma \left(\beta \left[\log \frac{p_\theta(\mathbf{x}_{0:T}^w \mid c)}{p_{\text{ref}}(\mathbf{x}_{0:T}^w \mid c)} - \log \frac{p_{\text{mix}}(\mathbf{x}_{0:T}^l \mid c)}{p_{\text{ref}}(\mathbf{x}_{0:T}^l \mid c)} \right] \right) \right], \quad (24)$$

714 where the mixture distribution is defined as

715
$$p_{\text{mix}}(\mathbf{x}_{0:T}^l \mid c) = \lambda p_\theta(\mathbf{x}_{0:T}^l \mid c) + (1 - \lambda) p_{\text{ref}}(\mathbf{x}_{0:T}^l \mid c), \quad \lambda \in (0, 1). \quad (25)$$

717 To obtain a computationally tractable form, we first decompose $\log p_\theta(\mathbf{x}_{0:T} \mid c)$ using the backward
 718 process of diffusion models:

720
$$\log p_\theta(\mathbf{x}_{0:T} \mid c) = \log p_\theta(\mathbf{x}_T \mid c) + \log \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1} \mid \mathbf{x}_t, c) \quad (26)$$

723
$$= \log p_\theta(\mathbf{x}_T \mid c) + \sum_{t=1}^T \log p_\theta(\mathbf{x}_{t-1} \mid \mathbf{x}_t, c) \quad (27)$$

726
$$= C - \sum_{t=1}^T \omega(t) \|\epsilon_\theta(x_t, c, t) - \epsilon\|_2^2, \quad (28)$$

730 where C is a parameter-independent constant.

731 Next, consider the mixed log-probability:

732
$$\begin{aligned} \log p_{\text{mix}}(\mathbf{x}_{0:T}^l \mid c) &= \log \left(\lambda p_\theta(\mathbf{x}_{0:T}^l \mid c) + (1 - \lambda) p_{\text{ref}}(\mathbf{x}_{0:T}^l \mid c) \right) \\ 734 &= C + \log \left(\lambda e^{\sum_{t=1}^T \log p_\theta(\mathbf{x}_{t-1}^l \mid \mathbf{x}_t, c)} + (1 - \lambda) e^{\sum_{t=1}^T \log p_{\text{ref}}(\mathbf{x}_{t-1}^l \mid \mathbf{x}_t, c)} \right). \end{aligned} \quad (29)$$

736 For brevity, define

738
$$a_t := \log p_\theta(\mathbf{x}_{t-1}^l \mid \mathbf{x}_t, c), \quad b_t := \log p_{\text{ref}}(\mathbf{x}_{t-1}^l \mid \mathbf{x}_t, c), \quad t = 1, \dots, T. \quad (30)$$

739 Then equation 29 becomes

741
$$\log p_{\text{mix}}(\mathbf{x}_{0:T}^l \mid c) = C + \log \left(\lambda e^{\sum_{t=1}^T a_t} + (1 - \lambda) e^{\sum_{t=1}^T b_t} \right). \quad (31)$$

743 Applying the discrete n -factor Hölder inequality yields

744
$$\lambda e^{\sum_{t=1}^T a_t} + (1 - \lambda) e^{\sum_{t=1}^T b_t} \leq \prod_{t=1}^T \left(\lambda e^{p_t a_t} + (1 - \lambda) e^{p_t b_t} \right)^{1/p_t}, \quad (32)$$

747 for exponents $p_1, \dots, p_T \geq 1$ with $\sum_{t=1}^T 1/p_t = 1$.

748 Taking logs and substituting back gives

750
$$\log p_{\text{mix}}(\mathbf{x}_{0:T}^l \mid c) \leq C + \sum_{t=1}^T \frac{1}{p_t} \log \left(\lambda e^{p_t a_t} + (1 - \lambda) e^{p_t b_t} \right). \quad (33)$$

753 Choosing $p_t = T$ for all t (so that $\sum_t 1/p_t = 1$) produces the uniform bound

754
$$\log p_{\text{mix}}(\mathbf{x}_{0:T}^l \mid c) \leq C + \frac{1}{T} \sum_{t=1}^T \log \left(\lambda e^{T \log p_\theta(\mathbf{x}_{t-1}^l \mid \mathbf{x}_t, c)} + (1 - \lambda) e^{T \log p_{\text{ref}}(\mathbf{x}_{t-1}^l \mid \mathbf{x}_t, c)} \right). \quad (34)$$

756 Table 6: **Comparison of Diff-SNPO with baseline methods on SD1.5 and SDXL backbones**
757 **on HPDv2.** All values are reported as mean \pm 95% confidence interval over 4 random seeds. Diff-
758 SNPO achieves the best scores across most human preference metrics, indicating superior alignment
759 and visual quality. For each metric, the top-performing method is **bolded**, while the second-best is
760 underlined.

762 Model	763 Method	764 HPSv2	765 Pick Score	766 Aesthetic Score	767 Image Reward
768 SD1.5	Baseline	26.52 ± 0.10	20.80 ± 0.06	5.3406 ± 0.17	-0.0521 ± 0.19
	Diff-DPO (Wallace et al., 2024)	26.85 ± 0.03	21.26 ± 0.02	5.4716 ± 0.14	0.1620 ± 0.08
	Diff-NPO (Wang et al., 2025)	27.18 ± 0.06	21.71 ± 0.05	5.5842 ± 0.07	0.2287 ± 0.14
	CHATS (Fu et al., 2025)	27.68 ± 0.07	21.45 ± 0.04	5.8605 ± 0.07	0.3699 ± 0.08
	Diff-BDPO (Ours)	27.24 ± 0.03	21.57 ± 0.03	5.5677 ± 0.06	0.3576 ± 0.02
	Diff-SNPO (Ours)	27.85 ± 0.06	22.86 ± 0.03	<u>5.7737 ± 0.06</u>	0.8093 ± 0.08
769 SDXL	Baseline	28.01 ± 0.15	22.80 ± 0.09	5.9887 ± 0.08	0.8781 ± 0.11
	Diff-DPO (Wallace et al., 2024)	28.55 ± 0.09	23.12 ± 0.01	6.0296 ± 0.05	1.0788 ± 0.09
	Diff-NPO (Wang et al., 2025)	28.87 ± 0.03	23.29 ± 0.05	6.0816 ± 0.01	1.1305 ± 0.04
	CHATS (Fu et al., 2025)	28.82 ± 0.07	22.85 ± 0.01	5.9868 ± 0.03	1.1165 ± 0.02
	Diff-BDPO (Ours)	28.62 ± 0.04	22.96 ± 0.10	5.9562 ± 0.09	1.0911 ± 0.10
	Diff-SNPO (Ours)	28.89 ± 0.09	23.34 ± 0.03	5.9876 ± 0.03	1.1460 ± 0.04

774 **Upper bound on the objective.** Substituting equation 34 and equation 28 into equation 24, we
775 obtain

$$776 \mathcal{L}_{\text{Diff-BDPO}}(\theta) \leq -\mathbb{E}_{(\mathbf{x}_{0:T}^w, \mathbf{x}_{0:T}^l, c) \sim \mathcal{D}} \left[\log \sigma \left(\mathbb{E}_{t \sim \mathcal{U}[1, T]} \left[-\beta m(x_t^w, c) - m_{\text{mix}}(x_t^l, c) \right] \right) \right] \quad (35)$$

$$777 \leq -\mathbb{E}_{(x_0^w, x_0^l, c) \sim \mathcal{D}, t \sim \mathcal{U}[1, T]} \left[\log \sigma \left(-\beta (m(x_t^w, c) - m_{\text{mix}}(x_t^l, c)) \right) \right] \\ 778 = \mathcal{L}_{\text{Diff-BDPO-UB}}(\theta), \quad (\text{by Jensen's inequality}). \quad (36)$$

783 Here the per-step terms are defined as

$$784 d_\theta(x_t, \epsilon, t, c) = T \omega(t) \|\epsilon - \epsilon_\theta(x_t, t, c)\|_2^2, \quad d_{\text{ref}}(x_t, \epsilon, t, c) = T \omega(t) \|\epsilon - \epsilon_{\text{ref}}(x_t, t, c)\|_2^2, \quad (37)$$

$$785 m(x_t, c) = d_\theta(x_t, \epsilon, t, c) - d_{\text{ref}}(x_t, \epsilon, t, c), \quad (38)$$

$$786 m_{\text{mix}}(x_t, c) = -\log \left(\lambda e^{-d_\theta(x_t, \epsilon, t, c)} + (1 - \lambda) e^{-d_{\text{ref}}(x_t, \epsilon, t, c)} \right) - d_{\text{ref}}(x_t, \epsilon, t, c). \quad (39)$$

789 In summary, by decomposing the diffusion likelihood and applying Hölder's inequality, we obtain
790 a tractable upper bound on the original Diffusion-BDPO objective, expressed in terms of per-step
791 denoising errors equation 37–equation 39.

793 A.3 ADDITIONAL QUANTITATIVE RESULTS

795 Table 6 presents the results on the HPDv2 test set, showing trends consistent with those observed
796 on the Pick-a-Pic benchmark. On the SD1.5 backbone, Diff-SNPO delivers a clear improvement,
797 raising HPSv2 to 22.86 and ImageReward to 0.81, underscoring its effectiveness in capturing hu-
798 man preferences. On the more advanced SDXL backbone, it continues to perform strongly, re-
799 maining competitive with the leading approach, Diff-NPO. Crucially, Diff-SNPO achieves this per-
800 formance with significantly lower computational cost and faster sampling, benefits enabled by its
801 single-network design. This efficiency advantage highlights its practicality, offering state-of-the-art
802 alignment quality while maintaining scalability and runtime efficiency.

803 A.4 ABLATION STUDY ON DIFFERENT ODE SOLVERS

805 In this section, we conduct an ablation study to assess the performance of Diff-SNPO across different
806 ODE solvers. To this end, we evaluate four widely used solvers: DDIM (Song et al., 2021), Euler
807 Discrete (Karras et al., 2022), UniPC (Zhao et al., 2023), and DPM Solver (Lu et al., 2022). The
808 results, presented in Table 7, show the performance of Diff-SNPO on the Pick-a-Pic v2 dataset with
809 SD1.5. From these results, we observe that the choice of solver has little impact on the performance
metrics. This outcome aligns with the theory behind Diffusion ODE solvers, where different solvers

810 Table 7: Performance comparison of different samplers across various reward metrics.
811

812 Sampler	813 Hpsv2 \uparrow	814 Aesthetic Score \uparrow
815 DDIM	816 27.23	5.6258
Euler Discrete	27.22	5.6460
UniPC	27.25	5.6468
Dpm Solver	27.25	5.6466

818
819 are alternative numerical methods for solving the same ODE system (Lu et al., 2022; Karras et al.,
820 2022). Consequently, with a sufficiently large number of function evaluations (NFE), all solvers
821 converge to the same image, leading to negligible differences in performance.
822

823 A.5 SAFETY ALIGNMENT RESULTS 824

825 Table 8: Comparison of IP values before and after finetuning on a CoProv2 dataset.
826

827 SafetyDPO	828 IP \downarrow
829 Baseline	0.4308
830 Diff-DPO	0.5109
831 Diff-NPO	0.4203
832 Diff-SNPO	0.4719
After Finetuning on CoProv2	
834 Diff-DPO	0.1713
835 Diff-NPO	0.1318
836 Diff-SNPO	0.1100

837
838 Table 8 presents the results of training and evaluating models on the CoProv2 (Wu et al., 2025) for
839 SD 1.5. This dataset contains 23,690 pairs of safe and unsafe images, spanning across 7 categories
840 ((Hate, Harassment, Violence, Self-Harm, Sexual, Shocking, Illegal). Evaluation was performed
841 using the Inappropriate Probability (IP) metric (Schramowski et al., 2023), which quantifies the
842 model’s ability to generate safe content when prompted with unsafe prompts.
843

844 From Table 8, we observe that models trained on the Pick-a-Pic v2 dataset exhibit weaker safety
845 performance. Specifically, their IP scores increase relative to the baseline after finetuning, indicating
846 a decline in safety. This is primarily due to the nature of the Pick-a-Pic v2 dataset itself, which
847 contains some unsafe images within its “win” samples. As a result, models trained on this dataset
848 are inadvertently exposed to unsafe content, causing them to generate unsafe outputs.
849

850 In contrast, when explicitly trained with a safety-oriented dataset, the safety performance improves
851 significantly. In particular, Diff-SNPO outperforms Diff-DPO, achieving a safety score of 0.11,
852 which is slightly better than Diff-NPO’s score of 0.13. This demonstrates that using a safety-focused
853 dataset, in combination with Diff-SNPO, can lead to improved alignment with safety objectives.
854

855 A.6 TOTAL TRAINING COMPUTE COMPARISONS 856

857 Table 9 shows the total GPU hours required by different baseline methods and Diff-SNPO for both
858 SD1.5 and SDXL using A6000 GPUs. The GPU hours for the baselines were calculated based on
859 the hyperparameters specified in the respective papers. From Table 9, it is clear that Diff-SNPO
860 outperforms its counterparts in computational efficiency, demonstrating its ability to achieve strong
861 performance with significantly lower resource requirements.
862

863

864
865

Table 9: Total GPU hours for different methods on SD 1.5 and SDXL using A6000 GPUs..

866
867
868
869
870
871
872
873
874
875
876

Model	Method	Total GPU Hours
SD 1.5	Diff-DPO	206
	CHATS	237
	Diff-NPO	2×206
	Diff-SNPO	79
SDXL	Diff-DPO	1592
	CHATS	2162
	Diff-NPO	2×1592
	Diff-SNPO	498

877
878

A.7 INFERENCE COMPUTATIONAL COST

879
880
881
882
883
884
885

Table 10 reports the inference throughput of Diff-SNPO and other negative preference optimization methods on SD1.5. The efficiency of the single-model design is particularly evident at inference: Diff-SNPO achieves the highest throughput among all methods, processing more than twice as many images per second as the dual-model baselines.

886
887
888
889

By contrast, CHATS and Diff-NPO suffer from slower inference due to the need for two sequential forward passes—one for each model—combined with additional passes required by their sampling schemes. These results underscore Diff-SNPO’s ability to perform fast, parallel inference without compromising alignment quality.

890
891
892
893
894

Overall, our findings demonstrate that Diff-SNPO effectively balances preference alignment with computational efficiency. Its single-model design reduces training cost while delivering faster inference, making it well-suited for large-scale training as well as real-time or resource-constrained deployment scenarios.

895
896

A.8 QUALITATIVE RESULTS

897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917

Figure 5 and Figure 6 present additional image samples generated by Diff-SNPO and various preference alignment methods on SD 1.5 and SDXL using prompts from HPDv2 test set respectively.

Table 10: Inference cost comparison.
Single-image inference cost with SD1.5 on a single A6000 GPU.

Method↓	Throughput (img/s) ↑
CHATS	0.11
DiffNPO	0.27
Diff-SNPO (Ours)	0.48

918	Prompt	SD1.5	Diff.-NPO	CHATS	Diff.-SNPO (Ours)
919	A small dog looking at a white plate holding donuts.				
920					
921					
922					
923					
924	A jellyfish sleeping in a space station pod.				
925					
926					
927					
928	A cat in a tutu dancing to Swan Lake.				
929					
930					
931					
932					
933					
934					
935					
936					
937					
938					
939					
940	Spider-Man holding a ginger cat.				
941					
942					
943					

Figure 5: Side-by-side comparison of images generated by related methods on HPDv2 using SD1.5.

947	Prompt	SDXL	Diff.-NPO	CHATS	Diff.-SNPO (Ours)
948	A lemon wearing a suit and tie, full body portrait.				
949					
950					
951					
952					
953					
954					
955					
956	Side-view blue-ice sneaker inspired by Spiderman created by Weta FX.				
957					
958					
959					
960					
961					
962					
963	2B from NieR Automata eating a bagel.				
964					
965					
966					
967					
968					
969					
970					
971	A spaceship in an empty landscape.				

Figure 6: Side-by-side comparison of images generated by related methods on HPDv2 using SDXL.